arXiv:1405.7741v2 [math.OC] 27 Jun 2014

Properties of Pseudocontractive Updates in Convex Optimization

Patrick W. Gallagher PATRICK.W.GALLAGHER@GMAIL.COM
UC San Diego, 9500 Gilman Dr., La Jolla, CA 92093

Zhuowen Tu ZTU@UCSD.EDU
UC San Diego, 9500 Gilman Dr., La Jolla, CA 92093

Abstract

Many convex optimization methods are conceived of and analyzed in a largely separate fashion. In contrast to this traditional separation, this manuscript points out and demonstrates the utility of an important but largely unremarked common thread running through many prominent optimization methods. In particular, we show that methods such as successive orthogonal projection, gradient descent, projected gradient descent, the proximal-point method, forward-backward splitting, the alternating direction method of multipliers, and under- or over-relaxed variants of the preceding all involve updates that are of a common type — namely, the updates satisfy a property known as pseudocontractivity. Moreover, since the property of pseudocontractivity is preserved under both composition and convex combination, updates constructed via these operations from pseudocontractive updates are themselves pseudocontractive. Having demonstrated that pseudocontractive updates are to be found in many optimization methods, we then provide a unified basic analysis of methods with pseudocontractive updates. Specifically, we prove a novel bound satisfied by the norm of the difference in iterates of pseudocontractive updates and we then use this bound to establish that the error criterion $\left\|x^N - Tx^N\right\|^2$ is $o(1/N)$ for any method involving pseudocontractive updates (where N is the number of iterations and T is the iteration operator).

Originally submitted January 30, 2014 to the 31^{st} International Conference on Machine Learning. Copyright 2014 by the author(s).

1. Introduction

Optimization methods are traditionally viewed and analyzed largely separately from one another. In this manuscript, we highlight a common feature, namely *pseudocontractivity*, of the iterative updates of many methods, and we show how recognizing this commonality can lead to a unified analysis and deeper understanding of a variety of prominent convex optimization methods.

We summarize our contributions as follows: (1) we highlight the essentially unremarked prevalence of pseudocontractive updates in many convex optimization methods; (2) we prove a novel result, Lemma 1, establishing a bound on the decrease of the norm of the displacement operator associated with any pseudocontractive operator; (3) we use the Lemma 1 result on monotone norm decrease to prove a novel result, Theorem 3, establishing that the error criterion $\left\|x^N - Tx^N\right\|^2$ is $o(1/N)$ for any method involving pseudocontractive updates (where N is the number of iterations and T is the iteration operator).

In order to establish the results above, we also bring together a variety of previously widely dispersed results; the combination of these results turns out to greatly enhance their collective utility and applicability.

1.1. Related Work

We believe that pseudocontractive operators provide a new and widely useful perspective in the context of machine learning (and largely unmentioned in optimization). Most standard optimization texts (e.g., (Bertsekas, 1999; Bertsekas et al., 2003; Nocedal & Wright, 2006; Boyd & Vandenberghe, 2004)) omit discussion of any of the related classes of operators that form the basis of our discussion; namely operators that are averaged, inverse strongly monotone, or pseudocontractive. We now briefly consider those references that do mention one (or, in some cases, two) of these operator classes.

In (Eckstein, 1989) the discussion covers (projected) gradient descent, the proximal-point method, forward-backward splitting, the alternating direction method of multipliers, and numerous other methods. The coverage does touch on the special case of 1-pseudocontractivity (i.e., firm nonexpansiveness) and on the special case of 1-inverse strong monotonicity; however, without the more general concepts of ν-pseudocontractivity and σ-inverse strong monotonicity (and the relationship between these concepts) there are significant limits to the results that can be obtained. One particularly notable limitation of a focus on firmly nonexpansive operators is that the class of firmly nonexpansive operators is not closed under composition.

(Vasin & Ageev, 1995) discusses both averaged operators and pseudocontractive operators, and also establishes the relationship between these classes; however, the general concept of inverse strong monotonicity is not present; in addition to the absence of the combined relationship between pseudocontractivity and inverse strong monotonicity, there is no unified convergence rate analysis of the type we discuss. On the other hand, (Bauschke & Combettes, 2011) discusses averagedness and inverse strong monotonicity (called there co-coercivity). Both averagedness and inverse strong monotonicity are discussed extensively and the latter is the central property on which most of the subsequent analysis is based. However, pseudocontractivity is essentially absent and convergence rates are not discussed.

(Golshtein & Tretyakov, 1996) discusses inverse strong monotonicity extensively in their discussion of gradient-type methods and in methods involving modified monotone mappings; however, the concepts of averagedness and pseudocontractivity are absent. The discussion in (Byrne, 2004) establishes that many iterative optimization methods have updates that are *averaged*; this work also includes the relationship between averaged and inverse strongly monotone. However, the concept of pseudocontractivity is absent. While the Krasnoselskii-Mann Theorem (Mann, 1953) is used to establish convergence of methods with averaged updates, there is no analysis of convergence rates.

Of more recent papers discussing proximal methods, such as (Combettes & Wajs, 2005; Combettes & Pesquet, 2011), the special case of firm nonexpansiveness appears, but not the more general cases of interest to us. The story is similar for other recent papers discussing the alternating direction method of multipliers, such as (Boyd et al., 2011; Parikh & Boyd, 2013; He & Yuan, 2012). This, again, limits the utility and breadth of applicability of the results obtained.

In the machine learning community, as in the optimization community, the classes of operators that we consider pass largely unremarked. (Lee et al., 2012) uses the proximal-point mapping and establishes the firm nonexpansiveness of some of the operations involved in their algorithm. The discussion in (Duchi & Singer, 2009; Singer & Duchi, 2009) squarely covers some of the methods that we touch on here; despite an apparent similarity in general outline, the analysis is pursued from a rather different perspective that is largely disjoint from our approach; in part this is due to a focus there on objective function value, in contrast to our focus on the sequence of iterates. One of the methods we discuss, the alternating direction method of multipliers, has recently been considered in stochastic and online variants (Ouyang et al., 2013; Wang & Banerjee, 2012), but our discussion here is restricted to the batch deterministic case.

1.2. Preliminaries

We restrict our attention to the finite dimensional case; this avoids situations in which the methods we discuss might attain weak convergence but not strong convergence (Güler, 1991) (since in the finite-dimensional case, the concepts of weak convergence and strong convergence coincide).

We use the notation Fix T to denote the fixed point set of the operator T; that is, Fix $T = \{x \in \mathbb{R}^n \mid T(x) = x\}$; we restrict our attention throughout to operators having nonempty fixed point sets. We use $\|\cdot\|$ to denote the standard 2-norm and $\langle \cdot, \cdot \rangle$ to denote the associated standard inner product. When convenient, we switch freely between operator notation and relation notation; for example, we might state the operator inclusion $y \in T(x)$ equivalently in relation form as $(x, y) \in T$.

With regard to terminology, we note that the property of inverse strong monotonicity (Rockafellar & Wets, 2004; Golshtein & Tretyakov, 1996) is elsewhere called co-coercivity (Combettes, 2009). On the other hand, the term pseudocontractive is sometimes used to mean a significantly different property (cf (Bertsekas & Tsitsiklis, 1989)) than our usage here.

2. Classes of Operators

We begin with a number of definitions.

Definition 1. An operator $T : \mathbb{R}^n \to \mathbb{R}^n$ is *nonexpansive* when

$$\|Tx - Ty\| \leq \|x - y\| \tag{1}$$

for all $x, y \in \mathbb{R}^n$.

Definition 2. For any operator $T : \mathbb{R}^n \to \mathbb{R}^n$, the associated *displacement operator* is $G = I - T$.

Definition 3. An operator $T : \mathbb{R}^n \to \mathbb{R}^n$ is *firmly nonexpansive* when

$$\|Tx - Ty\|^2 \leq \|x - y\|^2 - \|Gx - Gy\|^2, \tag{2}$$

for all $x, y \in \mathbb{R}^n$, where G is the displacement operator associated with T.

Definition 4. An operator $T : \mathbb{R}^n \to \mathbb{R}^n$ is ν-*pseudocontractive* when

$$\|Tx - Ty\|^2 \leq \|x - y\|^2 - \nu \|Gx - Gy\|^2 \quad (3)$$

for all $x, y \in \mathbb{R}^n$, where $\nu > 0$ and G is the displacement operator associated with T.

Note that firmly nonexpansive operators correspond to ν-pseudocontractive operators in the special case of $\nu = 1$.

Definition 5. An operator $T : \mathbb{R}^n \to \mathbb{R}^n$ is α-*averaged* when it is of the form

$$T = (1 - \alpha)I + \alpha N, \quad (4)$$

where N is a nonexpansive operator and $\alpha \in (0, 1)$.

We remark in passing that the use of the term averaged as defined above corresponds to the generalization from the special case in which the convex combination parameter $\alpha = 1/2$.

Definition 6. A point-to-set operator $T : \mathbb{R}^n \to 2^{\mathbb{R}^n}$ is *monotone* when

$$\langle u - v, x - y \rangle \geq 0 \quad (5)$$

for any $(x, u) \in T$, $(y, v) \in T$, where $x, y \in \mathbb{R}^n$.

Definition 7. An operator $T : \mathbb{R}^n \to \mathbb{R}^n$ is σ-*inverse strongly monotone* when

$$\langle Tx - Ty, x - y \rangle \geq \sigma \|Tx - Ty\|^2 \quad (6)$$

for all $x, y \in \mathbb{R}^n$, where $\sigma > 0$.

2.1. Pseudocontractive Updates: Prevalence and Common Convergence Rate Analysis

We discuss the details in later sections after necessary theoretical developments; for now we state the two main themes of this manuscript.

Proposition (Details in Section 3)**.** *The following prominent convex optimization methods have pseudocontractive updates:*

a) *Orthogonal projection onto a nonempty closed convex set is* 1*-pseudocontractive.*

b) *Over- or -under relaxed orthogonal projection methods are* $(\frac{2-\omega}{\omega})$*-pseudocontractive, where* $\omega \in (0, 2)$ *is the relaxation parameter.*

c) *For a convex function* f *with an* L*-Lipschitz gradient, the gradient descent iterative update operator* $I - \gamma \nabla f$ *is* $(\frac{2}{\gamma L} - 1)$*-pseudocontractive for any step factor* $\gamma \in (0, \frac{2}{L})$.

d) *The proximal-point mapping,* $\text{prox}_{f,\lambda}$*, associated with a convex function* f *is* 1*-pseudocontractive for any* $\lambda > 0$.

e) *The* ω*-relaxed proximal-point mapping,* prox_{ω}*, where* $\omega \in (0, 2)$*, associated with a convex function* f *is* $(\frac{2-\omega}{\omega})$*-pseudocontractive for any* $\lambda > 0$.

f) *Projected gradient descent involving a convex function* f *with an* L*-Lipschitz gradient is pseudocontractive with parameter* $(1 - \frac{\gamma L}{2})$*, where* $\gamma \in (0, \frac{2}{L})$ *is the gradient descent step factor.*

g) *Over- or -under-relaxed projected gradient descent is pseudocontractive.*

h) *Forward-backward splitting in which the smooth convex function* f *has an* L*-Lipschitz gradient is pseudocontractive with parameter* $(1 - \frac{\gamma L}{2})$*, where* $\gamma \in (0, \frac{2}{L})$ *is the gradient descent step factor.*

i) *The most basic version of the alternating direction method of multipliers update is* 1*-pseudocontractive.*

In addition to highlighting the common role played by pseudocontractive updates in all of the methods above, we also present a novel proof establishing that the error criterion $\left\|x^N - Tx^N\right\|^2$ is $o(1/N)$ for any method with a pseudocontractive update.

Theorem (Details in Section 4)**.** *The error criterion* $\left\|x^N - Tx^N\right\|^2$ *is* $o(1/N)$ *for any method with a pseudocontractive iteration operator.*

Having stated our primary results, we now begin introducing the background concepts that we will use in establishing these results.

2.2. Relationships

We note that the following statements are equivalent:

- T is $\left(\frac{1-\alpha}{\alpha}\right)$-pseudocontractive
- T is α-averaged
- $I - T$ is $\frac{1}{2\alpha}$-inverse strongly monotone

The relationship between pseudocontractive and averaged can be found in (Vasin & Ageev, 1995). The relationship between averaged and inverse strongly monotone can be found in (Byrne, 2004). The combination of all three of the concepts of pseudocontractivity, averagedness, and inverse strong monotonicity in the present manuscript provides the foundation for most of our results.

We also note an equivalence in a special case: T being 1-pseudocontractive is equivalent to T being 1-inverse strongly monotone; see Corollary 1.

The significance of these relationships is that they allow us to first establish properties in whatever form is most immediate and then to translate the properties to the most convenient form for subsequent use.

2.2.1. Over- or under-relaxation involving a firmly nonexpansive operator

Finally, we note that if we start with a firmly nonexpansive operator F, it is actually the case that the operator $T = (1-\omega)I + \omega F$ is averaged not only for $\omega \in (0,1)$, but for $\omega \in (0,2)$. This is because when F is firmly nonexpansive, its associated reflection[1] operator $R = 2F - I$ is nonexpansive (see Proposition 8).

We state this result on relaxation in the following Proposition for future reference:

Proposition 1. *For a firmly nonexpansive operator F, the operator $T = (1-\omega)I + \omega F$ is $(\frac{2-\omega}{\omega})$-pseudocontractive for $\omega \in (0,2)$.*

Proof. Denoting the reflection operator associated with F as $R = 2F - I$, we note that

$$(1-\omega)\, I + \omega F = \left(1 - \frac{\omega}{2}\right) I + \frac{\omega}{2} R. \tag{7}$$

Since Proposition 8 establishes that R is nonexpansive, the expression above means that the operator $T = (1-\omega)I + \omega F$ can be seen to be $\frac{\omega}{2}$-averaged for $\omega \in (0,2)$; this in turn corresponds to T being $(\frac{2-\omega}{\omega})$-pseudocontractive for $\omega \in (0,2)\square$.

The significance of the result above is that iterations involving under- or over-relaxation of a pseudocontractive update operator are also covered by our results. Over-relaxation, in which $\omega \in (1,2)$, is a broadly useful technique most commonly used to accelerate convergence of iterative methods in numerical linear algebra (Varga, 2009); the observation that over-relaxation of a firmly nonexpansive operator yields a pseudocontractive update is thus of some notable interest.

2.3. Closure Properties

The convex combination of nonexpansive operators is a nonexpansive operator. The composition of nonexpansive operators is a nonexpansive operator. The convex combination of pseudocontractive operators is a pseudocontractive operator. The composition of pseudocontractive operators is a pseudocontractive operator. These results can be found in (Vasin & Ageev, 1995); however, we can in fact be even more specific:

Theorem 1 ((Ogura & Yamada, 2002) Theorem 3)**.** *When T_1 is α_1-averaged, where $\alpha_1 \in [0,1)$, and T_2 is α_2-averaged, where $\alpha_2 \in [0,1)$, we have*

a) *for $t \in [0,1]$, the t-convex combination $T_{cc} = (1-t)T_1 + tT_2$ is α_{cc}-averaged, where $\alpha_{cc} = (1-t)\alpha_1 + t\alpha_2$,*

b) *the composition $T_{co} = T_1 T_2$ is α_{co}-averaged, where $\alpha_{co} = \frac{\alpha_1+\alpha_2-2\alpha_1\alpha_2}{1-\alpha_1\alpha_2}$.*

Using the relationship between averaged and pseudocontractive, we note that the convex combination T_{cc} is ν_{cc}-pseudocontractive, where $\nu_{\text{cc}} = \frac{1-\alpha_{\text{cc}}}{\alpha_{\text{cc}}}$; similarly, the composition T_{co} is ν_{co}-pseudocontractive, where $\nu_{\text{co}} = \frac{1-\alpha_1-\alpha_2+\alpha_1\alpha_2}{\alpha_1+\alpha_2-2\alpha_1\alpha_2}$.

Note that while this result establishes that both T_1T_2 and T_2T_1 are pseudocontractive, the respective fixed point sets need not coincide[2].

The significance of these closure properties is twofold: first, closure under composition and convex combination means that new algorithms can very naturally be constructed by using previous algorithms as building blocks; second, the resulting constructed algorithms will maintain the desirable properties that the individual building block updates possessed (and using Theorem 1 we can track the associated pseudocontractivity parameter values from the simple methods to the more complex combined methods).

3. Methods with Pseudocontractive Updates

In this section we demonstrate that pseudocontractive updates are widespread in many popular convex optimization methods.

3.1. Projection

Most discussions of projection methods only establish that projection is nonexpansive; the following known result demonstrates that projection is in fact 1-pseudocontractive (i.e., firmly nonexpansive).

Proposition 2. *Orthogonal projection $P_{\mathcal{C}}$ onto a nonempty closed convex set $\mathcal{C} \subset \mathbb{R}^n$ is 1-pseudocontractive. For convenience we show the equivalent result that projection is*

[1] The use of the term reflection for this operator is most immediately clear when we consider the case of projection onto a hyperplane H. Denoting projection onto the hyperplane H by P_H and denoting reflection across the hyperplane by R_H, we observe that $P_H = {}^1\!/_2 I + {}^1\!/_2 R_H$. The usage in the case above is a natural generalization.

[2] For example, consider the case of successive orthogonal projection onto two closed, convex, nonempty sets $\mathcal{A}$ and $\mathcal{B}$, with $\mathcal{A} \cap \mathcal{B} = \emptyset$. In this situation, $P_{\mathcal{A}}P_{\mathcal{B}}$ and $P_{\mathcal{B}}P_{\mathcal{A}}$ each have nonempty fixed point sets, but Fix $P_{\mathcal{A}}P_{\mathcal{B}}$ is the set of points in $\mathcal{A}$ that are at the minimal distance to $\mathcal{B}$, whereas Fix $P_{\mathcal{B}}P_{\mathcal{A}}$ is the set of points in $\mathcal{B}$ that are at the minimal distance to $\mathcal{A}$.

1-*inverse strongly monotone; that is,*

$$\langle P_{\mathcal{C}}x - P_{\mathcal{C}}y, x - y\rangle \geq \|P_{\mathcal{C}}x - P_{\mathcal{C}}y\|^2, \tag{8}$$

for all $x, y \in \mathbb{R}^n$.

Proof. The variational characterization of projection is: z is the projection of x onto $\mathcal{C}$ if and only if $\langle c - z, x - z\rangle \leq 0$ for every $c \in \mathcal{C}$ (see, e.g., (Ruszczynski, 2006)); we will express this as

$$\langle c - P_{\mathcal{C}}x, x - P_{\mathcal{C}}x\rangle \leq 0, \tag{9}$$

for every $c \in \mathcal{C}$. From the preceding characterization, we have, for any $x, y \in \mathbb{R}^n$

$$\langle P_{\mathcal{C}}x - P_{\mathcal{C}}y, y - P_{\mathcal{C}}y\rangle \leq 0 \tag{10}$$
$$\langle P_{\mathcal{C}}y - P_{\mathcal{C}}x, x - P_{\mathcal{C}}x\rangle \leq 0, \tag{11}$$

yielding the immediate conclusion

$$\|P_{\mathcal{C}}x - P_{\mathcal{C}}y\|^2 \leq \langle P_{\mathcal{C}}x - P_{\mathcal{C}}y, x - y\rangle, \tag{12}$$

for any $x, y \in \mathbb{R}^n$. The result immediately above establishes that $P_{\mathcal{C}}$ is 1-inverse strongly monotone, so that Corollary 1 implies that $P_{\mathcal{C}}$ is 1-pseudocontractive. □

3.2. Over- or Under-relaxed Projection

Using the notation P for orthogonal projection onto the nonempty closed convex set $\mathcal{C}$, we introduce the ω-relaxed projection operator $P_\omega = (1-\omega)I + \omega P$, where $\omega \in (0, 2)$. Since we just established that P is 1-pseudocontractive, we conclude that the associated reflection operator $2P - I$ is nonexpansive, and so from our previous reasoning in Equation 7 we observe that P_ω is $(\frac{2-\omega}{\omega})$-pseudocontractive.

We note in passing that reasoning similar to the reasoning in the previous two subsections can also be used to establish that cyclic subgradient projection (CSP) methods (in both their relaxed and unrelaxed forms) (Censor & Zenios, 1997) involve pseudocontractive updates.

3.3. Gradient Descent

Typical discussions of gradient descent do not make any reference to operator theory, or, more specifically, to pseudocontractivity; the known result below establishes conditions under which the gradient descent iterative update is pseudocontractive. The explicit form of the gradient descent iterative update is

$$x^{k+1} = [I - \gamma\nabla f]\left(x^k\right). \tag{13}$$

Proposition 3. *For a convex function f with an L-Lipschitz gradient, the gradient descent iterative update operator $I - \gamma\nabla f$ is $(\frac{2}{\gamma L} - 1)$-pseudocontractive for any step factor $\gamma \in (0, \frac{2}{L})$.*

Proof. We have assumed that ∇f is L-Lipschitz, which implies that $\frac{1}{L}\nabla f$ is nonexpansive. Theorem 6.9 from Chapter 1 of (Golshtein & Tretyakov, 1996) establishes that $\frac{1}{L}\nabla f$ nonexpansive implies that $\frac{1}{L}\nabla f$ is 1-pseudocontractive.

Corollary 1 establishes that $\frac{1}{L}\nabla f$ being 1-pseudocontractive is equivalent to $\frac{1}{L}\nabla f$ being 1-inverse strongly monotone. Lemma 4 establishes that $\frac{1}{L}\nabla f$ being 1-inverse strongly monotone implies that $\gamma\nabla f$ is $\frac{1}{\gamma L}$-inverse strongly monotone, since $\gamma L > 0$.

When an operator T is exactly $1/2$-inverse strongly monotone, $G = I - T$ is nonexpansive. When T is σ-inverse strongly monotone for $\sigma \in (1/2, +\infty)$, $G = I - T$ is $(2\sigma - 1)$-pseudocontractive. To ensure $I - \gamma\nabla f$ is pseudocontractive, we must thus ensure that $\gamma\nabla f$ is σ-inverse strongly monotone with $\sigma \in (1/2, +\infty)$; equivalently, we must ensure that $(2\sigma - 1) > 0$.

$\gamma\nabla f$ being $\frac{1}{\gamma L}$-inverse strongly monotone corresponds to $I - \gamma\nabla f$ being $\left(\frac{2}{\gamma L} - 1\right)$-pseudocontractive. The pseudocontractivity parameter is required to be strictly positive; in order to ensure that $\left(\frac{2}{\gamma L} - 1\right) > 0$, we require $\gamma \in \left(0, \frac{2}{L}\right)$, since $\left(\frac{2}{\gamma L} - 1\right) > 0$ requires $\frac{2}{L} > \gamma$. □

3.4. Proximal-Point Method

As is the case with projection, most discussions of proximal-point methods only establish that the proximal-point mapping is nonexpansive; the following known result demonstrates that the proximal-point mapping is in fact 1-pseudocontractive (i.e., firmly nonexpansive). We note that the proximal-point mapping associated with f can be written in either of the two equivalent forms

$$\text{prox}_{f,\lambda}(x) = \underset{z\in\mathbb{R}^n}{\text{argmin}} \left\{ f(z) + \frac{1}{2\lambda}\|z - x\|^2 \right\} \tag{14}$$
$$\text{prox}_{f,\lambda}(x) = (I + \lambda\partial f)^{-1}(x), \tag{15}$$

where $\lambda > 0$; see, e.g., (Rockafellar, 1970). The mapping $(I + \lambda\partial f)^{-1}$ is called the λ-resolvent of the subdifferential mapping ∂f.

Proposition 4. *For a convex function f, the associated proximal-point update is 1-pseudocontractive, for any $\lambda > 0$.*

Proof. Consider $x^+ = \text{prox}_{f,\lambda}(x)$ and $y^+ = \text{prox}_{f,\lambda}(y)$. Lemma 3 tells us that

$$\left(x^+, x - x^+\right) \in \lambda\partial f \tag{16}$$
$$\left(y^+, y - y^+\right) \in \lambda\partial f. \tag{17}$$

The subdifferential mapping of a convex function is monotone (Rockafellar, 1970); thus $\lambda\partial f$, where $\lambda > 0$, is also

monotone. From this we have

$$\left\langle x^+ - y^+, \left(x - x^+\right) - \left(y - y^+\right)\right\rangle \geq 0 \tag{18}$$

$$\left\langle x^+ - y^+, (x - y) - \left(x^+ - y^+\right)\right\rangle \geq 0 \tag{19}$$

$$\left\langle x^+ - y^+, x - y\right\rangle \geq \left\|x^+ - y^+\right\|^2. \tag{20}$$

This establishes that $\text{prox}_{f,\lambda}$ is 1-inverse strongly monotone; applying Corollary 1 then establishes that $\text{prox}_{f,\lambda}$ is 1-pseudocontractive for any $\lambda > 0$. □

Although we do not go into further detail, we note here that analogous arguments apply to any algorithm involving the resolvent operator associated with a maximal monotone mapping (cf (Eckstein, 1989)).

3.5. Relaxed proximal-point Method

Lightening our notation and for the moment taking prox to denote the λ-proximal-point mapping associated with the convex function f, we introduce the ω-relaxed proximal-point iteration operator $\text{prox}_\omega = (1-\omega)I + \omega\text{prox}$, where $\omega \in (0, 2)$. Since we have established that prox is 1-pseudocontractive, we conclude that the associated reflection operator $2\text{prox} - I$ is nonexpansive, and so from our previous reasoning in Equation 7 we observe that prox_ω is $(\frac{2-\omega}{\omega})$-pseudocontractive.

3.6. Projected Gradient Descent

The iterative update for projected gradient descent involving a convex function f with an L-Lipschitz gradient and a nonempty closed convex constraint set $\mathcal{C}$ corresponds to the composition $P_\mathcal{C} \circ [I - \gamma\nabla f]$. We have established that the projection $P_\mathcal{C}$ is 1-pseudocontractive and that the gradient descent update $I - \gamma\nabla f$ is $(\frac{2}{\gamma L} - 1)$-pseudocontractive, where $\gamma \in (0, \frac{2}{L})$. Since the composition of pseudocontractive updates is pseudocontractive (Vasin & Ageev, 1995), we have thus established that projected gradient descent is pseudocontractive. We also note that under- or over-relaxation of the projection portion of the update still yields pseudocontractivity.

We can use Theorem 1b to be more specific. Basic projected gradient descent is the composition of the gradient descent step, which is $\frac{\gamma L}{2}$-averaged, and the projection step, which is $1/2$-averaged; the composition thus is $(1 - \frac{\gamma L}{2})$-pseudocontractive. Under- or over-relaxed projected gradient descent has the $1/2$-averaged projection replaced with ω-relaxed projection, which we have seen is $(\frac{2-\omega}{\omega})$-pseudocontractive.

3.7. Forward-Backward Splitting

Consider the optimization problem

$$\underset{x\in\mathbb{R}^n}{\text{minimize}}\ f(x) + g(x), \tag{21}$$

where f and g are each convex and f has an L-Lipschitz gradient. The explicit statement of the forward-backward splitting method is

$$x^{k+1} = \text{prox}_{g,\lambda} \circ [I - \gamma\nabla f](x^k) \tag{22}$$

When used as a gradient descent step followed by a proximal step, so long as the gradient descent step factor is $\gamma \in (0, \frac{2}{L})$, we have the composition of pseudocontractive operators, which is in turn pseudocontractive.

We can again use Theorem 1b to be more specific. Basic forward-backward splitting is the composition of the gradient descent step, which is $\frac{\gamma L}{2}$-averaged, and the proximal step, which is $1/2$-averaged; the composition thus is $(1 - \frac{\gamma L}{2})$-pseudocontractive. When we use under- or over-relaxation of the proximal step in forward-backward splitting, the $1/2$-averaged proximal step is replaced with an ω-relaxed generalized proximal step, which we have seen is $(\frac{2-\omega}{\omega})$-pseudocontractive.

3.8. Alternating Direction Method of Multipliers

Consider the optimization problem

$$\underset{x\in\mathbb{R}^n}{\text{minimize}}\ f(x) + g(x), \tag{23}$$

where f and g are each convex. Using the notation $R_f = 2\text{prox}_{f,\lambda} - I$ and $R_g = 2\text{prox}_{g,\lambda} - I$, we can express (Boyd, 2011) the most basic version of the alternating direction method of multipliers (ADMM) update in terms of operators as

$$x^{k+1} = [1/2 I + 1/2 R_g R_f](x^k). \tag{24}$$

Proposition 4 establishes that $\text{prox}_{f,\lambda}$ and $\text{prox}_{g,\lambda}$ are each 1-pseudocontractive, for any $\lambda > 0$. Proposition 8 then establishes that the respective reflection operators R_f and R_g are each nonexpansive. The composition of nonexpansive operators is nonexpansive, so $R_g R_f$ is nonexpansive. We thus observe that the most basic version of the ADMM update is $1/2$-averaged, and thus is 1-pseudocontractive. We also immediately observe that the more general convex-combination ADMM-like update

$$x^{k+1} = [(1-\alpha)I + \alpha R_g R_f](x^k), \tag{25}$$

where $\alpha \in (0, 1)$, will be $(\frac{1-\alpha}{\alpha})$-pseudocontractive.

4. Convergence Rate of Methods with ν-Pseudocontractive Updates

The worst-case convergence rate is well-known for each of the methods that we discuss in Section 3. The standard approach generally requires separate arguments to establish this rate for each method; in contrast, by leveraging the fact that the updates in each of these methods are pseudocontractive (and making use of the relationships between pseudocontractive, averaged, and inverse strongly monotone) we are able to use one argument for all of these methods.

Consider a method for which the iteration operator T is ν-pseudocontractive. We will establish

Lemma 1. *For a ν-pseudocontractive iteration operator T, the norm of the displacement operator $G = I - T$ satisfies the expression*

$$\left\|Gx^{k+1}\right\|^2 \leq \left\|Gx^k\right\|^2 - \nu \left\|Gx^k - Gx^{k+1}\right\|^2 \tag{26}$$

for $k \in \{0, 1, \ldots\}$.

Proof. From the relationship between pseudocontractivity and inverse strong monotonicity, we note that when T is ν-pseudocontractive, we have that $G = I - T$ is $\frac{1+\nu}{2}$-inverse strongly monotone. That is, we have

$$\langle Gx - Gy, x - y\rangle \geq \left(\frac{1+\nu}{2}\right)\|Gx - Gy\|^2, \tag{27}$$

for all $x, y \in \mathbb{R}^n$. Thus,

$$\left\langle Gx^k - Gx^{k+1}, x^k - x^{k+1}\right\rangle \geq \frac{1+\nu}{2}\left\|Gx^k - Gx^{k+1}\right\|^2 \tag{28}$$

$$\left\|Gx^k\right\|^2 - \left\langle Gx^{k+1}, Gx^k\right\rangle \geq \frac{1+\nu}{2}\left\|Gx^k - Gx^{k+1}\right\|^2, \tag{29}$$

which yields

$$\left\|Gx^{k+1}\right\|^2 \leq \left\|Gx^k\right\|^2 - \nu\left\|Gx^k - Gx^{k+1}\right\|^2, \tag{30}$$

where we have used $Gx^k = x^k - x^{k+1}$ and $2\left\langle Gx^{k+1}, Gx^k\right\rangle = \left\|Gx^{k+1}\right\|^2 + \left\|Gx^k\right\|^2 - \left\|Gx^{k+1} - Gx^k\right\|^2$. □

Using Lemma 1, we can now show that

Theorem 2. *Any method with a ν-pseudocontractive iteration operator has error criterion satisfying $\left\|x^N - Tx^N\right\|^2 \leq \frac{1}{(N+1)\nu}\left\|x^0 - x^\star\right\|^2$ at iteration N.*

Proof. An alternate expression of ν-pseudocontractivity with one argument, $x^\star$, coming from the fixed point set is

$$\nu\left\|Gx^k\right\|^2 \leq \left\|x^k - x^\star\right\|^2 - \left\|x^{k+1} - x^\star\right\|^2. \tag{31}$$

Summing over all the iterations from 0 to N yields

$$\sum_{k=0}^{N}\left\|Gx^k\right\|^2 \leq \frac{1}{\nu}\left\|x^0 - x^\star\right\|^2. \tag{32}$$

Lemma 1 tells us that $\left\|Gx^{k+1}\right\|^2 \leq \left\|Gx^k\right\|^2$ for each k, so we conclude that

$$\left\|Gx^N\right\|^2 \leq \frac{1}{(N+1)\nu}\left\|x^0 - x^\star\right\|^2. \square \tag{33}$$

We can establish an even stronger result by leveraging a useful technical Lemma from (Dong, 2013).

Lemma 2 ((Dong, 2013))**.** *Consider three sequences of strictly positive numbers $\{\alpha_k\}, \{\beta_k\}, \{\gamma_k\}$. When these sequences are such that $\{\beta_k\}$ is unsummable, $\{\gamma_k\}$ is nonincreasing, and the relationship $\alpha_{k+1}^2 \leq \alpha_k^2 - \beta_k\gamma_k$ holds for each $k \in \{0, 1, \ldots\}$, then it will also be the case that there exists another sequence $\{\varepsilon_k\}$ such that*

- $\varepsilon_k \in (\alpha_k, \alpha_0)$,
- $\lim_{k\to+\infty}\varepsilon_k = \lim_{k\to+\infty}\alpha_k$,
- $\gamma_k \cdot \left[\sum_{i=0}^{k}\beta_i\right] \leq 2\cdot\alpha_0\cdot\varepsilon_k$.

We now use this Lemma to establish

Theorem 3. *For any method with a ν-pseudocontractive iteration operator, the error criterion $\left\|x^N - Tx^N\right\|^2$ is $o(1/N)$.*

Proof. We will apply Lemma 2 with $\alpha_k = \left\|x^k - x^\star\right\|$, $\beta_k = \nu$, and $\gamma_k = \left\|Gx^k\right\|^2$. An infinite sequence of constants is clearly unsummable and Lemma 1 establishes that $\{\left\|Gx^k\right\|^2\}$ is nonincreasing.

Moreover, the definition of ν-pseudocontractivity when one argument, $x^\star$, comes from the fixed point set provides us with the necessary relationship

$$\left\|x^{k+1} - x^\star\right\|^2 \leq \left\|x^k - x^\star\right\|^2 - \nu\left\|Gx^k\right\|^2 \tag{34}$$

for each $k \in \{0, 1, \ldots\}$.

Thus, Lemma 2 tells us that there exists a sequence $\{\varepsilon_k\}$ such that

- $\varepsilon_k \in \left(\left\|x^k - x^\star\right\|, \left\|x^0 - x^\star\right\|\right)$,
- $\lim_{k\to+\infty}\varepsilon_k = \lim_{k\to+\infty}\left\|x^k - x^\star\right\|$,
- $\left\|Gx^k\right\|^2 \cdot \left[\sum_{i=0}^{k}\nu\right] \leq 2\cdot\left\|x^0 - x^\star\right\|\cdot\varepsilon_k$.

Taking limsup on both sides of the preceding inequality and noting that $\lim_{k\to+\infty}\varepsilon_k$ exists and is equal to $\lim_{k\to+\infty}\left\|x^k-x^\star\right\|$ we have

$$\limsup_{k\to\infty}\left\{\left\|Gx^k\right\|^2\left[\sum_{i=0}^{k}\nu\right]\right\}\leq\lim_{k\to\infty}\left\{2\left\|x^0-x^\star\right\|\varepsilon_k\right\} \tag{35}$$

$$\limsup_{k\to\infty}\left\{\left\|Gx^k\right\|^2\left[\nu\left(k+1\right)\right]\right\}\leq 2\left\|x^0-x^\star\right\|\lim_{k\to\infty}\left\|x^k-x^\star\right\| \tag{36}$$

$$\limsup_{k\to\infty}\left\{\left\|Gx^k\right\|^2\left[\nu\left(k+1\right)\right]\right\}\leq 0, \tag{37}$$

where we have used $\lim_{k\to+\infty}\left\|x^k-x^\star\right\|=0$ (for our finite dimensional setting, this follows because, e.g., the Krasnoselskii-Mann Theorem (Mann, 1953) establishes weak convergence of methods with averaged updates [and thus with pseudocontractive updates] and weak convergence coincides with strong convergence in finite dimension).

The final inequality above establishes the desired result: in terms of iteration N, we observe that $\left\|Gx^N\right\|^2$ is $o(1/N)$.□

5. Discussion

Pseudocontractivity is a property largely overlooked by the machine learning community. This paper takes a first step toward showing the benefits of recognizing and using pseudocontractivity. We have demonstrated the broad prevalence of pseudocontractive updates in existing optimization methods. We also emphasize that optimization methods made up of updates derived from compositions or convex combinations of pseudocontractive operators will also be pseudocontractive — this provides a perspective on the convergence behavior of combined methods that appears to be entirely absent from classical analyses. We provide a demonstration of a unified analysis that applies to all pseudocontractive methods.

With an eye to following these ideas from theory to practice, we are extending the analysis in this paper to the settings of inexact and stochastic pseudocontractive updates. Recent computational experience suggests that allowing inexact updates can have provide a significant boost to the practical efficiency of many optimization methods. We also anticipate that a Nesterov-style estimate sequence approach will provide a basis for accelerated versions of any pseudocontractive method. We believe that all of these characteristics provide strong motivation for researchers to examine optimization methods for unlooked-for pseudocontractivity.

A. Appendix

The following known results are necessary for establishing the results elsewhere in the manuscript.

Proposition 5. *An operator F is 1-pseudocontractive if and only if the displacement operator $G=I-F$ is 1-pseudocontractive.*

Proof. Immediate from the definition of 1-pseudocontractive. □

Proposition 6. *An operator F is 1-inverse strongly monotone if and only if the displacement operator $G=I-F$ is 1-inverse strongly monotone.*

Proof. The relevant expressions for either direction are

$$\|Gx-Gy\|^2\leq\langle Gx-Gy,x-y\rangle \tag{38}$$

$$\|(x-y)-(Fx-Fy)\|^2\leq\langle x-y,x-y\rangle-\langle Fx-Fy,x-y\rangle \tag{39}$$

$$\|Fx-Fy\|^2\leq\langle x-y,Fx-Fy\rangle\,. \text{ □} \tag{40}$$

Proposition 7. *An operator F is 1-pseudocontractive if and only if the displacement mapping $G=I-F$ is 1-inverse strongly monotone.*

Proof. The relevant expressions for either direction are

$$\|Tx-Ty\|^2\leq\|x-y\|^2-\|Gx-Gy\|^2 \tag{41}$$

$$\|Tx-Ty\|^2+\|[x-y]-[Tx-Ty]\|^2\leq\|x-y\|^2 \tag{42}$$

$$\|Tx-Ty\|^2\leq\langle x-y,Tx-Ty\rangle\,. \text{ □} \tag{43}$$

Corollary 1. *An operator F is 1-pseudocontractive if and only if F is 1-inverse strongly monotone.*

Proof. Proposition 7 shows that F is 1-pseudocontractive if and only if $G=I-F$ is 1-inverse strongly monotone. Proposition 6 shows that $G=I-F$ is 1-inverse strongly monotone if and only if F is 1-inverse strongly monotone. Together these establish the result. □

Proposition 8 ((Goebel & Kirk, 1990), Theorem 12.1)**.** *For a firmly nonexpansive operator F, the reflection operator $R=2F-I$ is nonexpansive.*

Proof.

$$\|Rx-Ry\|^2=4\|Fx-Fy\|^2+\|x-y\|^2-4\langle Fx-Fy,x-y\rangle \tag{44}$$

$$\leq\|x-y\|^2, \tag{45}$$

for all $x,y\in\mathbb{R}^n$, where we have used the fact that F firmly nonexpansive implies that F is 1-inverse strongly monotone; that is, that $\langle Fx-Fy,x-y\rangle\geq\|Fx-Fy\|^2$, for all $x,y\in\mathbb{R}^n$. □

Lemma 3. *For a convex function f, the proximal-point $x^+ = prox_{f,\lambda}(x)$ satisfies the inclusion*

$$x - x^+ \in \lambda \partial f(x^+), \tag{46}$$

where $\lambda > 0$.

Proof. The resolvent statement of the proximal-point mapping is

$$\text{prox}_{f,\lambda}(x) = (I + \lambda \partial f)^{-1}(x). \tag{47}$$

From $x^+ = \text{prox}_{f,\lambda}(x)$, we thus have

$$\left(x, x^+\right) \in (I + \lambda \partial f)^{-1} \tag{48}$$

$$\left(x^+, x\right) \in I + \lambda \partial f \tag{49}$$

$$x - x^+ \in \lambda \partial f\left(x^+\right), \tag{50}$$

where we have freely shifted our view from relation to operator as convenient. □

Lemma 4. *T being σ-inverse strongly monotone implies that, for $\gamma > 0$, the scaled operator γT is $\frac{\sigma}{\gamma}$-inverse strongly monotone.*

Proof. Immediate from the definition of inverse strongly monotone. □

References

Bauschke, Heinz H and Combettes, Patrick L. *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2011.

Bertsekas, Dimitri P. *Nonlinear programming*. Athena Scientific, 2nd edition, 1999.

Bertsekas, Dimitri P. and Tsitsiklis, J.N. *Parallel and distributed computation*. Old Tappan, NJ (USA); Prentice Hall Inc., 1989.

Bertsekas, Dimitri P., Nedic, Angelia, and Ozdaglar, Asuman E. *Convex Analysis and Optimization*. Athena Scientific, 2003.

Boyd, S. P. Monotone Operator Splitting Methods. *Unpublished note*, 2011.

Boyd, S. P. and Vandenberghe, L. *Convex optimization*. Cambridge Univ Pr, 2004.

Boyd, S. P., Parikh, N., Chu, E., Peleato, B., and Eckstein, Jonathan. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.

Byrne, Charles. A unified treatment of some iterative algorithms in signal processing and image reconstruction. *Inverse Problems*, 20(1):103, 2004.

Censor, Y. and Zenios, S.A. *Parallel optimization: Theory, algorithms, and applications*. Oxford University Press, USA, 1997.

Combettes, Patrick L. Fejér monotonicity in convex optimization. In *Encyclopedia of Optimization*, pp. 1016–1024. Springer, 2009.

Combettes, Patrick L and Pesquet, Jean-Christophe. Proximal splitting methods in signal processing. In *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pp. 185–212. Springer, 2011.

Combettes, P.L. and Wajs, V.R. Signal Recovery by Proximal Forward-Backward Splitting. *Multiscale Modeling & Simulation*, 4:1168, 2005.

Dong, Yunda. The Proximal Point Algorithm Revisited. *Journal of Optimization Theory and Applications*, pp. 1–12, 2013.

Duchi, John and Singer, Yoram. Efficient online and batch learning using forward backward splitting. *The Journal of Machine Learning Research*, 10:2899–2934, 2009.

Eckstein, Jonathan. *Splitting methods for monotone operators with applications to parallel optimization*. PhD thesis, Massachusetts Institute of Technology, 1989.

Goebel, Kazimierz and Kirk, William A. *Topics in metric fixed point theory*, volume 28. Cambridge University Press, 1990.

Golshtein, E.G. and Tretyakov, NV. *Modified Lagrangians and monotone maps in optimization*. Wiley-Interscience, 1996.

Güler, Osman. On the convergence of the proximal point algorithm for convex minimization. *SIAM Journal on Control and Optimization*, 29(2):403–419, 1991.

He, Bingsheng and Yuan, Xiaoming. On the O(1/n) Convergence Rate of the Douglas-Rachford Alternating Direction Method. *SIAM Journal on Numerical Analysis*, 50(2):700–709, 2012.

Lee, Jason, Sun, Yuekai, and Saunders, Michael. Proximal Newton-type methods for convex optimization. In *Advances in Neural Information Processing Systems*, pp. 836–844, 2012.

Mann, W Robert. Mean value methods in iteration. *Proceedings of the American Mathematical Society*, 4(3): 506–510, 1953.

Nocedal, J. and Wright, S.J. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, New York, 2nd edition, 2006.

Ogura, Nobuhiko and Yamada, Isao. Non-strictly convex minimization over the fixed point set of an asymptotically shrinking nonexpansive mapping. *Numerical functional analysis and optimization*, 23(1):113–138, 2002.

Ouyang, Hua, He, Niao, Tran, Long, and Gray, Alexander G. Stochastic alternating direction method of multipliers. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pp. 80–88, 2013.

Parikh, Neal and Boyd, Stephen. Proximal algorithms. *Foundations and Trends in Optimization*, pp. 1–96, 2013.

Rockafellar, Ralph Tyrrell. *Convex analysis*. Princeton, NJ: Princeton University Press, XVIII, 451 pp, 1970.

Rockafellar, Ralph Tyrrell and Wets, Roger J-B. *Variational analysis*. Springer, 2004. ISBN 3540627723.

Ruszczynski, A.P. *Nonlinear optimization*. Princeton University Press, 2006.

Singer, Yoram and Duchi, John C. Efficient learning using forward-backward splitting. In *Advances in Neural Information Processing Systems*, pp. 495–503, 2009.

Varga, Richard S. *Matrix iterative analysis*, volume 27. Springer, 2009.

Vasin, Vladimir V and Ageev, Aleksandr L. *Ill-posed problems with a priori information*, volume 3 of *Inverse and Ill-posed Problems Series*. VSP, 1995.

Wang, Huahua and Banerjee, Arindam. Online Alternating Direction Method. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pp. 1119–1126, 2012.